Security Now! #872 - 05-24-22

Dis-CONTI-nued: The End of Conti?

This week on Security Now!

This week we'll start by following-up on Microsoft's patch Tuesday Active Directory domain controller mess. We're going to look at several instances of the Clearview AI facial recognition system making news and at the systems which fell during last week's Vancouver Pwn2Own competition. We cover some welcome news from the U.S. Department of Justice and some disturbing news about a relatively simple and obvious hack against popular Bluetooth-link smart locks. We have some closing the loop feedback from our listeners, including a look at what's going on with the Voyager 1 space probe, and another interesting look into the looming impact of quantum crypto. Then we finish by sharing an in-depth examination of the surprisingly deliberately orchestrated shutdown of the Conti ransomware operation.

Security News

Emergency mid-cycle update for Active Directory

Last week we noted that the previous week's patch Tuesday had been overly eventful with admins reporting that the month's patch roll-up had adversely affected the operation of some of their enterprise's critical authentication. The malfunctioning patch was intended to close an actively abused security vulnerability. And while there were some pre-patch workarounds, most admins were choosing to simply roll back their systems by removing the entire May Day patch bundle.

So, last week I wondered out loud whether Microsoft would attempt to fix the fix (and to hopefully test it more thoroughly this time) or whether they would prefer to wait until June's patch Tuesday. One problem with waiting is that June begins next Wednesday, which places June's patch Tuesday as late in month as possible — nearly in the middle, on the 14th.

In the past, Microsoft's patching logic has been inscrutable. Sometimes they'll wait half a year before patching something, other times they appear to be in a hurry. Fortunately, for Active Directory users, the latter was the case here with Microsoft fixing the trouble in about a week and releasing a mid-cycle emergency patch last Thursday. However, presumably since it was only domain controller servers that were adversely affected, the emergency update is not being made widely available through the regular Windows Update channel. Enterprising users will need to get it from Microsoft's Update catalog by going there directly. And enterprise users using Windows Server Update Services and Microsoft's Endpoint Configuration Manager can import the update into those offerings.

Since CISA was forced to pull the requirement of patching this flaw from their catalog of known exploited vulnerabilities due to May's debacle, we can assume that it will be returning. In any event, it's obvious that these days, updating is everything.

Clearview AI -vs- {Illinois, Australia, Canada and the United Kingdom}

We've talked about ClearView AI on several occasions. Recall that they're the company which provides what have turned out to be quite controversial facial recognition services to, apparently, anyone who wants them. The controversy surrounds the perceived privacy invasion created when an army of both visible and invisible cameras might be scanning and cataloging the real world identities of everyone who crosses through their fields of vision. 30 years ago we lacked the ability to automate anything like this "at scale" because it all comes down to economics. We didn't have inexpensive cameras, communications bandwidth, storage or computation. Today we have a seemingly endless supply of all of that for next to nothing in cost. So what could not be done in the recent past is now feasible on a mass scale. During a recent interview, ClearView AI's CEO, Hoan Ton-That explained that he intends to have 100 billion images in their database within a year, up from around 20 billion currently. That’s nearly 13 photos for each person on earth.

The potential for the abuse of this technology is breathtaking due to its sweeping scope. One of the times we have talked about these guys was when the ACLU, the American Civil Liberties Union, and others brought a suit against ClearView AI in Illinois’ thanks to Illinois' Biometric Information Privacy Act (BIPA) of 2008. Thanks to BIPA, any entity wishing to collect biometric

identifying information — such as a photograph — for any Illinois citizen must first obtain their explicit consent before doing so. Since ClearView AI scans online image sources, consent is never requested nor received.

Earlier this month, ClearView AI agreed to a settlement in that lawsuit brought by the ACLU that would largely prohibit sale of services to private companies and people in the United States, as well as to law enforcement throughout Illinois over the next five years. However, ClearView AI's CEO said that the settlement would not change anything on a material level because they could continue to sell to any government clients elsewhere.

At the same time, ClearView AI is increasingly facing pressure from various nation's privacy regulators who are pushing Clearview to remove data from its systems. Both Australia and Canada last year ordered the company to delete information on their residents.

And now, last week, the U.K. government announced that they are levying a fine of more than £7.5 million pound sterling (around $9.4 million USD) against Clearview AI, ordered it to stop collecting information about U.K. residents and to delete all of the data that it has already obtained, from its database. The U.K.'s Information Commissioner John Edwards said, in a press release: “The company not only enables identification of those people, but effectively monitors their behavior and offers it as a commercial service. That is unacceptable. People expect that their personal information will be respected, regardless of where in the world their data is being used. That is why global companies need international enforcement.” Edwards added that he would be meeting with European regulators in Brussels later this week to “collaborate to tackle global privacy harms.”

These recent moves by the Information Commissioner’s Office (ICO) follow a joint investigation by the agency and the Office of the Australian Information Commissioner that launched in July of 2020 and was completed last November. A provisional notice from ICO to Clearview that month warned the company to stop processing and delete U.K. resident data—as well as suggesting a substantially larger £17 million (pound sterling) fine.

In reply, Clearview's founder and chief executive Hoan Ton-That said he was “deeply disappointed” that the U.K. data authority “misinterpreted” his company’s technology and intentions. Uh huh. He said: “I would welcome the opportunity to engage in conversation with leaders and lawmakers so the true value of this technology which has proven so essential to law enforcement can continue to make communities safe.” he said.

Clearview AI in Ukraine

At the same time, there is no question that facial recognition technology can be quite useful. It's ClearView AI's facial recognition technology that has been allowing the Ukrainian government to identify both their own citizen casualties of the war with Russia, as well as the Russian soldier casualties who have been left behind by their comrades. Thanks to ClearView AI's knowledge of Russian citizens, Ukraine has been able to mount a potentially powerful political counteroffensive by identifying these Russian war casualties, notifying their families of their death, and offering to allow them to travel to Ukraine to reclaim their fallen family member. Some 400 Ukranian investigators are also using Clearview AI's technology to build much more airtight war crimes cases.

The Record recently interviewed Clearview AI's founder and CEO, Hoan Ton-That, whom I referred to before. I think that the details revealed by that interview are worth sharing:

How did Clearview come to play such a major role in Ukraine?

Hoan Ton-That: *When the war started, it was really shocking to me and a lot of members of our team to see especially the video footage of women and children suffering and it made me think, how can we help? Originally, I would see photos of captured Russian soldiers, and I realized with that kind of photo quality, our facial recognition technology could be helpful. So, I reached out to a lot of people on our advisory board to ask them, do you know anyone in the Ukrainian government, and one person — Lee Wolosky —he's on the National Security Council under three presidents, he quickly said yes. We thought that it could be helpful in identifying deceased soldiers and track misinformation.*

How could it be useful to track misinformation?

HT-T: *You'd see a lot of things on social media saying this is a captured Russian soldier, but you might see people on the other side saying actually that's a paid actor and here's the name of the paid actor. So with this technology, the level of accuracy can be used to identify if someone is who they say they are.*

Who in Ukraine has logins to Clearview AI?

HT-T: *It's in six different agencies, including the National Police of Ukraine. Then the people on the ground would be using it. So once we gave a demo, we'd give them training on how to use facial recognition technology responsibly. So part of the training is that they would send photos of unidentified people and run it through the system and we could show them, 'Hey, this is how you verify who it is.' So for example, if they have a tattoo that matches online, there's a very high chance it's the same person.*

You inserted this technology into a war zone, which presents a lot more problems than having a police department in the United States use it. How are you accounting for that?

HT-T: *You want to be careful to make sure that they really know how to use it properly, and so there are all these scenarios that we want to make sure don't happen. For example, what if someone takes a photo of someone and says, 'Hey, I think you're a Russian traitor.' And then they detain them. And it's all incorrect based on incorrect information. So we'd never want something like that to happen. As long as these investigations are done by trained investigators, this can be a very helpful tool. If it was used by everybody, I think that's when problems happen.*

What's the most surprising use you've seen in Ukraine?

HT-T: *War crimes investigations. We were talking to people in the National Police of Ukraine and others where they have video footage from surveillance cameras. Smartphones are more prevalent, so there's a higher chance of something being recorded. Now with this technology, if it's easy to identify someone, I think people are going to think twice about war crimes. So that*

was a surprise to me.

So this is a subscription service, and you say that gives you more control if someone is misusing it… how does Clearview AI work?

HT-T: *We vet every person to make sure they're a government official who is using it. There's also two-factor authentication, so they still have to verify their device before they log in. Once they have an account, there's an administrator for each agency. So, the administrator can see who is conducting searches and what reason they're conducting the search. There is an intake form that requires a case number and a crime type before conducting a search. So people when they're on-boarded and they're learning about the software, they know that their searches can be audited because we want to make sure they're using it for the right kind of stuff. Because it's a Cloud service, we have the ability to revoke access. If there's any egregious abuse of the technology, you want to make sure that we have the ability to take it away.*

The IT Army appears to be using it. In a video, the group demonstrated the use of a facial recognition program that appears to resemble Clearview AI. This is a volunteer hacking force, so how is it that the Ukraine IT Army appears to be using Clearview AI?

HT-T: *All I can say is that everyone we've on-boarded is a government official. We haven't onboard anyone in the IT Army directly. Everyone we talk to and on-board, we give them proper training on its usage. The speculation that the IT Army is running a Clearview AI search does not match any information we have on this matter. Clearview AI is intended for use in Ukraine by law enforcement and government officials.*

Did maybe somebody give a username and password to somebody who's in the IT Army?

HT-T: *It's possible that someone shared a screenshot or shared how it worked, but we want to make sure that whatever the usage of the technology is — say it is to identify someone deceased — that is done in a way that is positive. The policy of the National Police and all our users is to tell the family members in a humane way.*

Had it occurred to you that the IT Army would use this technology to notify families of dead soldiers as a propaganda tool?

HT-T: *I talked to some of the officials [in the Russian government] and I said, 'Look, is this something you knew about? Is that your procedure for doing this?' Then said that's not our official procedure. And they assured me that's not what they want to have happen either. Again, it is war time. Tensions are really high. Those things can happen.*

If I thought it would be used in a really bad way, then I don't think I'd give access to them. We think that just getting the information out in a humane way is the most important thing. What we can control as Clearview is giving access to the right people. So for example, we don't give access to the Russians or anything like that and we make sure Ukraine is trained as appropriately as possible.

Have you revoked any access related to Ukraine because you thought it wasn't being used properly?

HT-T: *No, not at this time, but the administrators of these agencies in Ukraine have the ability to do so. They can go in and audit the searches, remove access to an account and give access as they deem appropriate. Clearview AI would only revoke access to an agency if there was an egregious amount of abuse in that agency. Until something really escalates to that level, we haven't revoked any access.*

The NSO Group is an Israeli company that makes surveillance software that can be remotely implanted in smartphones. It has come under heavy criticism for its tech being used by authoritarian governments to spy on citizens. With your facial recognition technology, how do you avoid the NSO trap?

HT-T: *I think NSO is a very different kind of technology than what we do. We are searching public information from the internet. So it is really just like Google. If you can type someone's name and the word LinkedIn into Google and you find their photo, then Clearview can basically do the same thing, but it's search by photo. We also work with law enforcement. NSO is different because it's breaking into phones, very private data. Also, when they sell their software, they don't have the ability to take it back if they sell it to the wrong actor or government. Whereas Clearviews software is deployed in the cloud. If we ever find anything totally egregious or abusive, we have the ability to revoke access. We also want to be aligned with democratic countries, making sure that this is a technology that can be used responsibly and across democratic nations.*

Can you imagine a scenario, years and years from now when everyone has this capability that it would be in like VR glasses or built into a phone?

HT: *I can imagine like augmented reality is an interesting where it could be deployed on military bases so in Afghanistan, they had a situation where when they were pulling out at a checkpoint, terrorists could blow up people. They are looking up close at the ID's. To verify someone at a distance in a hands-free way, I think that's a very positive kind of use case.*

It's obviously trivial to invent synthetic positive use cases. It's at least as easy, probably much easier actually, to imagine quite privacy-intrusive use cases.

The logic we often hear being applied by those who object to the use of pervasive facial recognition is the assumption of privacy. They suggest that someone walking in a public space can be seen by everyone, so there's no assumption of privacy. But when a meta-tagged photo is taken among friends in a private setting and posted online to celebrate, is it reasonable for those images to be vacuumed up, identified, cataloged and made available for searches, globally? The argument is made that posting the photo online creates implied consent. But is that for any use whatsoever, forever? So we have another new problem created by astonishingly inexpensive technology. As a society we're going to have to figure out what's important and what we want and how we want these things to work.

Pwn2Own Vancouver 2022

Last Wednesday, Thursday and Friday, the 18th through the 20th, was the 15th anniversary of Pwn2Own held in Vancouver. Rather than enumerate the victories of many brilliant hackers who we don't know and many of their names which I cannot begin to pronounce correctly, I'm going to focus upon the products which we all know and which fell to their best efforts:

On Wednesday, the various hacker individuals and teams demonstrated an improper configuration against Microsoft Teams, an OOB Read and OOB Write to achieve privilege escalation on Oracle's Virtualbox, a 3-bug chain of injection, misconfiguration and sandbox escape against Microsoft Teams, prototype pollution and improper input validation against Firefox, an out-of-bounds write escalation of privilege on Windows 11, 2 bugs on Ubuntu Desktop - an Out-of-Bounds Write (OOBW) and Use-After-Free (UAF), a zero-click exploit of 2 bugs (injection and arbitrary file write) on Microsoft Teams, an out-of-band write on Apple Safari, a Use-After-Free elevation of privilege on Windows 11 and a Use-After-Free exploit on Ubuntu Desktop.

On the second day the hackers and some teams demonstrated 2 unique bugs (Double-Free & OOBW) with collision on a known sandbox escape on a Tesla Model 3 Infotainment System. There was also a Use After Free bug leading to elevation of privilege on Ubuntu Desktop, an improper access control bug leading to elevation of privilege on Windows 11, and a Use After Free bug leading to elevation of privilege on Ubuntu Desktop.

And on the third and final day as saw an escalation of privilege via Integer Overflow on Windows 11, a Use-After-Free exploit on Ubuntu Desktop, an Elevation of Privilege via Improper Access Control on Windows 11 and an Elevation of Privilege via Use-After-Free on Windows 11.

So the big targets and losers were Microsoft Teams, Windows 11 and Ubuntu Desktop, with VirtualBox, Firefox, Safari and Tesla's Infotainment system each taking one hit apiece. Overall, for this 15th anniversary event, there were a total of 21 exploitation attempts, three which failed, from 17 contestants with Trend Micro and the Zero Day Initiative awarding a total of $1,155,000 in PwnToOwn prize money.

The DoJ takes a welcome step back

Under the heading of “Sometimes they get it right” we have last Thursday's welcome news that the U.S. Department of Justice (our DOJ) has revised its policy on how federal prosecutors should charge violations under the Computer Fraud and Abuse Act, creating an explicit carve out exemption covering “good-faith” security research.

What this means is that the U.S. government is altering how vigorously it enforces a central cybercrime law that security researchers, civil liberty advocates and others have long-argued is overly broad. Amen!

Under the change to enforcement of the CFAA which became law back in 1986, the DoJ will amend its charging policy to explicitly discourage going after so-called “good faith,” or ethical, security researchers. The DoJ's Deputy Attorney General Lisa Monaco said in a statement accompanying the revamped policy: “Computer security research is a key driver of improved

cybersecurity. The department has never been interested in prosecuting good-faith computer security research as a crime, and today's announcement promotes cybersecurity by providing clarity for good-faith security researchers who root out vulnerabilities for the common good."

Federal prosecutors who seek to bring charges under CFAA must first consult with the Computer Crime and Intellectual Property unit inside DOJ's Criminal Division. If that office recommends going forward with charges, prosecutors must inform Monaco's team and even then may need special permission to proceed.

However, it's still not all clear sailing. Ethical researchers who scour for and discover software vulnerabilities could still face prosecution under existing **state** laws or be sued in civil court.

This updated guidance comes a little over a year after the Supreme Court ruled in a major CFAA case that the 1986 law does not apply when an authorized user utilizes data in improper ways. In that case, the court said a Georgia police officer did not violate the hacking law when he took money from an acquaintance to search a license plate database. In other words, someone paid a police officer to access the license plate database to which the officer had authorized access even though the purpose of that specific instance of access was unwarranted. The officer was sued under the Computer Fraud and Abuse Act and the Supreme Court said "nope." The DoJ explained that the CFAA should only apply in instances when an outside hacker or authorized user actually breaks into a secure portion of an organization's network. So, clearly, the instance of a police officer using this authorized access for an unauthorized purpose would not be a CFAA violation.

As expected, this updated policy was broadly welcomed among both federal cybersecurity officials and the researcher community. Jen Easterly, the director of CISA tweeted: "Huge news—well done, Team DOJ!" And Chris Vickery, a prominent cyber researcher, tweeted that the new guidance will "hopefully improve the lives of people (like me) who fear retaliation for trying to do the right thing."

So, as I said, "sometimes they get it right." We still have a ways to go, since local government and private parties who take offense at the discovery of their own cyber failings will still be able to somewhat unjustifiably exact their revenge. But perhaps a competent defense attorney will be able to point to the changes which have just been made at the federal level as a basis for applying the same reasonable protections more locally.

Sometimes, unlocking can be too convenient

So you have a fancy car, residential or office door lock which operates as follows:

Your smartphone containing the lock's matching app is in your pocket. You approach your car, your home's or your office's front door and just touch the lock. The lock's capacitive sensor senses your touch. So it sends out a Bluetooth Low Energy (BLE) ping to query for any nearby Kwikset Kevo smart lock apps. The app in the phone in your pocket receives the BLE ping and responds. So they engage in a super-crypto triple-scoop post-Quantum impossible-to-crack handshake negotiation, and, now satisfied that one of its authorized owners is nearby, the lock disengages to allow entry. And all is happy in mudville.

https://www.weiserlock.com/en/kevo/smart-lock

But now, we have a different scenario:

A bad guy team wants to gain entry into that smartlock-protected car, home or office. So one of them arranges to place a BLE relay near the user's phone. Perhaps in an adjacent office cubicle, next to their locker while they're working out, or in the coffee break room which they frequent. Wherever... it doesn't matter.

The other member of the team waits some discrete distance away from the car, home or office that's about to be breached. When the first team member messages to the second that the unsuspecting user's smartphone is within range of their remote Bluetooth relay, the second member of the team simply walks up to the locked car, home or office and touches its lock.

And exactly as before, the lock sends out a BLE discovery ping which the BLE relay forward to its matching endpoint. That distant endpoint simply and blindly rebroadcasts the ping, which the user's smartphone picks up and acknowledges. It replies and its reply is similarly forward back to the BLE relay which is now positioned near the lock. So once again they engage in the most unbreakable bazillion qubit quantum entangled crypto that the world has ever hosted, until the lock becomes satisfied that only the user's smartphone can possibly be at the other end of the link... so it disengages its lock to admit the individual whom it presumes is authorized to gain access.

Unfortunately, this is not fiction (except, in this example, for my overkill use of deeply entangled quantum crypto) — the point being, this oh-so-simple Bluetooth Low Energy relay attack entirely defeats any system's crypto, no matter whether pre- or post-quantum. And as I said, this is not fiction.

Last September, after successfully executing exactly this scenario in the field, the U.K's NCC Group notified several smart lock makers that they had a problem. Their systems were vulnerable to simple Bluetooth Low Energy relay attacks. And the NCC group notified the lock makers that they would eventually be publishing the news of this, which they did, last week.

The accompanying security advisory simply explains: *"An attacker within BLE signal range of a smartphone or key fob authorized to unlock a Kevo smart lock can conduct a relay attack to unlock the lock over long distances."*

And the problem is, there is no obvious in-band way of detecting and preventing this abuse.

Which is to say that BLE doesn't offer robust endpoint proximity protection. Various out-of-band solutions have been considered, one being for the smartphone to use GPS to ascertain whether it's physically near the lock it has been paired with. But concerns over GPS availability and speed have been a concern. And local GPS jamming and spoofing technology is available. Another out-of-band solution considered was to have the app monitoring its owner's physical movement and to disable any lock negotiation if the smartphone was motionless immediately before the negotiation began since that would never be expected in the system's normal use case. But any scenario where the user would be moving — even if many miles away — would defeat such anti-spoofing protection. So GPS would seem to be the best, if still imperfect, solution.

But the nature of the problem suggests that the use of BLE, convenient though it is, was not the best idea in the first place. If our Smartphones were equipped with radios which incorporated reliable time-of-flight measuring capabilities then it would be possible to obtain spoof-proof data **and** direct physical proximity assurance. But today's Smartphones are not yet quite that smart.

Closing The Loop

Brian Phillips / @dooq69

> *Hi Steve. Quick one... Don't know if you've covered this already but do you have an alternative for Unlock Origin on Safari? Recently got myself a MacBook Pro & now want to use Safari over Chrome. Thanks*

Ad Block Plus : https://adblockplus.org/ad-blocker-safari

Dave Badia / @DavB_F3chestnut

> *Can you give a shout out on SN that Voyager 1 is still happily trotting along after so many years? And I just read it goes into "safe mode" if things get out of whack so the programmers on earth can fix it!*

What's interesting is that Dave sent this last week, either before or more likely because things just started to go wonky with Voyager 1. Now, Voyager 1 has a Twitter account. (Of course it does.) Though being 14.5 billion miles away, with light requiring 21 hours, 34 minutes and 38 seconds to traverse the distance, NASA's public relations people have been posting on Voyager's behalf. Anyway, last Wednesday, "Voyager" posted on Twitter:

This made me curious about what was going on, so I dug around and found some commentary which I've lightly edited for the podcast. It reads:

NASA Voyager 1 space probe was launched 45 years ago and continues its journey as the first-ever human-made object to leave the vicinity of our solar system. It's heading out there to study the outer heliosphere and the interstellar medium. The iconic probe has sent hugely important data back to NASA since it launch on September 5th, 1977. But now, the new strange data sent by the Voyager from the edge of the Solar System has left scientists shocked since, until now, there have been no significant errors reported by the probe!

[And Leo, I have a somewhat chilling theory for what might be going on, but I'll finish this before sharing that…]

Since the Voyager 1 data is of critical importance, the engineering team is trying to solve the puzzle of mysterious data now being sent by the space probe. NASA said in a statement: "The interstellar explorer is operating normally, receiving and executing commands from Earth, along with gathering and returning science data. But readouts from the probe's Attitude Articulation and Control System (AACS) do not reflect what's actually happening onboard." But what does this data received back on earth actually mean? NASA says an antenna attached to Voyager, which is pointed at Earth to send data back, appears to be working but is sending back invalid data. The AACS controls the 45-year-old spacecraft's orientation. Among other tasks, it keeps Voyager 1's high-gain antenna pointed precisely at Earth, enabling it to send data home and to receive new instructions. All signs suggest the AACS is still working, but the telemetry data it's returning is invalid. For instance, the data may appear to be randomly generated, or does not reflect any possible state the AACS could be in. Thankfully, the issue

with the NASA Voyager 1 has not triggered the probe's onboard fault protection system, the system which is designed to keep the spacecraft in its "safe mode" which maintains a state where only essential operations are carried out, while giving engineers time to analyze and diagnose the issue. And Voyager 1's signals have not weakened, which suggests the high-gain antenna remains in its prescribed orientation to Earth.

Okay, so Leo... Nothing man made, and certainly nothing sending back telemetry from such a distance has ever been this far from Earth and its sun. What if Voyager is approaching the maximum radius supported by the simulation, and that's why its reality is beginning to break down and why it appears to be returning random data that makes no sense? Everything else we see out there appears to be just inbound radiation. It's not interactive. Voyager is... or at least it was.

It'll be interesting to see what happens, and whether perhaps the same thing happens to Voyager 2 once it gets around the same distance away. (Voyager is currently about 2.4 billion miles closer to Earth.)

https://voyager.jpl.nasa.gov/mission/status/

Reminiscent of the Mars rovers Spirit and Opportunity, both Voyagers have performed astonishingly well. Remember how the Gilligan's Island adventure started off as a three-hour tour? Similarly, the two Voyager spacecraft (and please, no one say "Veeger") were originally intended to study only Jupiter and Saturn, their moons, and Saturn's rings. And for that two-planet mission, they were built to last just five years... that was nearly 45 years ago. Their 45th anniversary of launch will be coming up this August and September.

And after their initial successes, NASA's engineers doubled the missions' objectives to include two more giant planets, Uranus and Neptune. So now, between the two, the spacecraft have explored four planets, 48 moons, a host of planetary magnetic fields and rings.

The Voyagers are each powered by a trio of radioisotope thermoelectric generators. These RTGs convert the heat generated by the radioactive decay of plutonium-238 into electricity. And while they are super-reliable and have no moving parts, that decay itself decays over time. Right now, the Voyagers' power is gradually dropping at a rate of 4 watts per year. This has been limiting the total number of systems that can be powered up at once and the Voyager mission team has been turning off non-essential equipment to reserve power. So far, however, none of Voyager's science instruments have been powered down. And the latest of the continually extending goals has been to keep both Voyagers running for at least another three years, beyond 2025.

Unless, of course, they spontaneously disappear on the other side of the simulation boundary.

Tom Feller / @ThomasJFeller

Hi Steve, I came across an interesting article I ran across on quantum computing: https://www.forbes.com/sites/arthurherman/2021/06/07/q-day-is-coming-sooner-than-we-think/ "Q-Day Is Coming Sooner Than We Think"

As we've discussed before, Forbes has had a somewhat spotty record for their presentation of technical subject matter. But assuming that the editors don't butcher it, the content's veracity will be a function of its author. In this case the article's author is a guy named Arthur Herman whose bio tells us that he's a Senior Fellow at the Hudson Institute, Director of the Quantum Alliance Initiative, and co-author of "Risking Apocalypse: Quantum Computers and the U.S. Power Grid."

Arthur is very clearly of the opinion that we should be much more worried than we currently are. Having read what he wrote, and even though I'm far from able to render any opinion on the subject, I know that our listeners will find this fascinating. Once again, I've edited what Forbes published so that it better fits our higher-end audience:

"Q-Day" is the term some experts use to describe when large-scale quantum computers are able to factor the large prime numbers that underlie our public encryption systems, such as the ones that are supposed to protect our bank accounts, financial markets, and most vital infrastructure. That's a feat that's all but impossible for even the fastest supercomputers but which the unique features of quantum computers, using the physics of superpositioning and entanglement, will be able to deliver.

There's a growing consensus that this quantum threat is real; there's no agreement how long it will take before a quantum computer has the 4000 or so stable qubits it will need to meet the requirements of Shor's algorithm for cracking those encryption systems.

For example, it would take a classical computer 300 trillion years to crack an RSA 2048-bit encryption key. A quantum computer can do the same job in just ten seconds with 4099 stable qubits — but getting to that number is the main problem quantum computer engineers face, since the stability or coherence of qubits lasts for only microseconds. Today's most entangled computer, Google's Bristlecone, has just 72 stable qubits.

Nonetheless, I have been arguing for the past four years, including in this column, that Q-Day is likely to come sooner than even quantum scientists predict, and that the time to get ready to protect our vulnerable data and networks is now. Others prefer to procrastinate, citing experts who say such a threat is at least a decade or more away. The fact that the National Institute of Standards and Technology (NIST) won't have its quantum-resistant algorithm standards ready until 2024, and expects the rollout to take another 5 to 15 years, has helped to encourage complacency disguised as confidence.

But new developments in quantum science suggest that this complacency may be misplaced. So-called quantum annealers like the one Canada-based D-Wave Systems, Inc. uses, are able to calculate the lowest energy level between the qubits' different states of entanglement, which equates to the optimal solution. These machines have proven their worth in solving optimization problems that usually stump classical computers, as I explained in a previous column.

Not surprisingly, scientists have been quietly finding ways to turn factorization into an optimization problem, thus bypassing Shor's algorithm, the paradigm for discussing quantum decryption since the 1990's. In 2019, scientific papers emerged that showed how to do this, including factoring integers using "noisy" qubits, i.e. swarms of quantum bits that aren't perfectly entangled the way a large-scale computer requires. In other words, no longer requiring stable qubits.

One was authored by Chinese scientists who found a way to factor a large number using only 89 noisy qubits. They then showed it's possible to factor an RSA 768-bit private key, which is the current factorization record accomplished using classical computers, with 147,454 noisy qubits — a fraction of the millions of noisy qubits a large quantum computer would need to reach the 4000 stable qubit threshold, and within reach of the architecture of an annealer like D-Wave Systems.

Also in 2019, a pair of Google researchers and the Royal Institute of Technology at Stockholm published a paper showing how to crack 2028-bit RSA integers in 8 hours using 20 million noisy qubits. Given the fact that in 2012 scientists speculated that it would take 1 billion qubits to perform this feat, it will likely not be long before researchers show they can get there with a lot fewer than 20 million qubits.

And sure enough, in 2020 three Chinese researchers found a way to use the D-Wave quantum computer to factor large integers, completely bypassing Shor's algorithm. "Thus," they concluded, "post-quantum cryptography should consider further the potential of the D-Wave quantum computer for deciphering the RSA cryptosystem in the future."

In effect, these researchers found a way to turn decryption using quantum technology into a straightforward process on a timeline much shorter than ten years: perhaps four to five years is more likely.

This was what Chinese scientists are openly publishing. We don't know what's happening behind the scenes, but we can bet if there's a short cut to achieve what a large-scale quantum computer can do using annealing technology, their military and intelligence services will want to find out.

All of this should change the timetable for Q-Day and for our strategic calculations. Not only is quantum-based decryption coming our way sooner, but thanks to annealing, this code-cracking feature will be more accessible to other machines than the hugely expensive large-scale computers Google, Microsoft, and others are working on — which, in turn, puts the capability within reach of small-state or even non-state actors.

Why gamble with the quantum future? Annealing technology makes becoming "quantum ready" more important, and getting started now, more imperative than ever.

https://www.nature.com/articles/s41598-020-62802-5.pdf

Dis-CONTI-nued: The End of “Conti”

Last Thursday, Advanced Intel’s (https://www.advintel.io/) Yelisey Boguslavskiy (@y_advintel) Tweeted: *“Today the official website of Conti #Ransomware was shut down, marking the end of this notorious crime group; it is truly a historic day in the #intelligence community!”*

And the day after, last Friday, Advanced Intel published their report about what happened. There’s so much more to it, that I felt certain our listeners would find the details fascinating.

https://www.advintel.io/post/discontinued-the-end-of-conti-s-brand-marks-new-chapter-for-cybercrime-landscape

Advanced Intel’s report is titled: *“DisCONTInued: The End of Conti’s Brand Marks New Chapter For Cybercrime Landscape”* ... and the top of their report teases:

> *From the negotiations site, chatrooms, messengers to servers and proxy hosts - the Conti brand, not the organization itself, is shutting down. However, this does not mean that the threat actors themselves are retiring.*

Okay, what does it mean? Advanced Intel rushed out their report which contained some typos, misspellings and grammatical awkwardness. In order to share it for this podcast I’ve cleaned it up a bit, but otherwise it remains what they wrote. And what they wrote is worth sharing:

> On May 19, 2022, the admin panel of the Conti ransomware gang's official website, Conti News, was shut down. The negotiations service site was also down, while the rest of the infrastructure: from chatrooms to messengers, and from servers to proxy hosts was going through a massive reset.
>
> Conti News - a shame blog - is the last beacon of the group’s public operation where victim data was published. It also served as a media tool that Conti used for their endless public statements (one of which led to the gang's downfall).
>
> This publicity function of the blog is still technically active (and this activity, as shown below, is highly strategized). At the time of this publication - May 20, 2022, Conti was even uploading anti-Americanist hate speech claiming the USA to be “a cancer on the body of the earth”. This, however, only manifests that the website became an empty shell. At the same time, the crucial operational function of Conti News which was to upload new data in order to intimidate victims to pay is defunct, as all the infrastructure related to negotiations, data uploads, and hosting of stolen data was shut down.

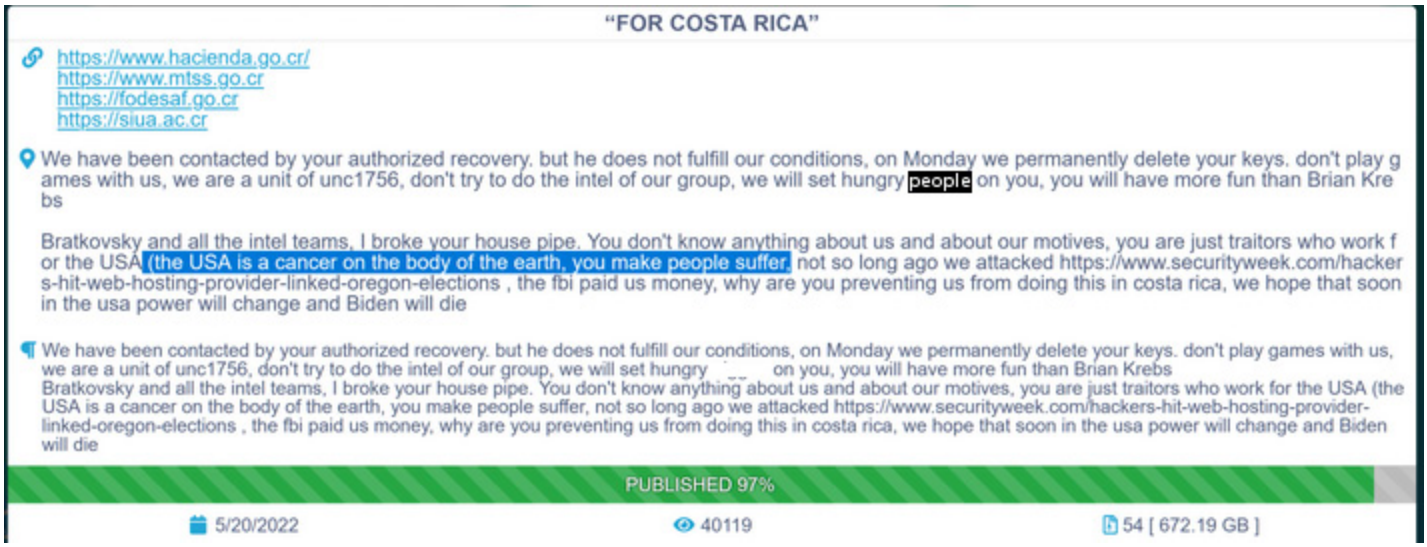

The Advanced Intel guys note that: *"The message published today is strikingly different from previous Conti's political statements written in properly edited English. The extremely low quality of writing also suggests that even the public side of the Conti blog is not treated seriously by the leadership."*

And this shutdown highlights a simple truth that has been evident for the Conti leadership since early Spring 2022 - the group can no longer sufficiently support and obtain extortion. The blog's key and only valid purpose is to leak new datasets, and this operation is now gone.

This was not a spontaneous decision, instead, it was a calculated move, signs of which were evident since late April. Two weeks ago, on May 6, AdvIntel explained that the Conti brand, and not the organization itself, was in the process of the final shutdown. As of May 19, 2022, our exclusive source intelligence confirms that today is Conti's official date of death.

In this retrospective analysis, we will not only take an in-depth look into the reasons behind the Conti shutdown but perhaps most importantly, assess and project the future of a new threat landscape that is already on the horizon. But first, we need to review how Conti prepared for its own demise, and how this group, notable for its sophistry, continued to utilize information warfare techniques to orchestrate the shutdown until its final days, in order to ensure the legacy of its surviving members:

Shutting down ransomware's iconic criminal brand is a long and complicated venture. A notorious and prolific threat group cannot simply turn off its servers, only to pop back up the following week with a new name and logo design. Even a whisper of novel threat group activity following the announcement of Conti's demise would likely spark immediate accusations of poorly executed identity theft. At best, immediate comparisons between the two would permanently leave the new group in Conti's ghostly shadow: the collective that fell and the one which emerged. [And I'll just note that we've seen and commented on exactly this previously.]

REvil, DarkSide, and countless other collectives attempted the disappearing act; the simple approach failed miserably. As what was one of the dominant ransomware groups active at the time, Conti realized that an element of "performativity" would need to be involved. Where other groups had been attempting a grand stunt with smoke and mirrors, Conti would try a sleight of hand.

Conti would not be itself without its project frontman - an individual operating under the alias "reshaev" aka "cybergangster". Besides being a talented coder (they were behind the original Ryuk payload), this person was an outstanding organizer. It was "reshaev" who set the foundation for Conti's dominance in the cybercrime business by creating an organizational system based on skill, teamwork, clear business processes, hierarchy, and clear foresight.

It is not surprising that "reshaev" was the first who saw Conti's structural challenges: due to the group's public allegiance to Russia in the first days of the Russian invasion into Ukraine, **Conti was not able to be paid**. Since February, almost no payments were given to the group, while Conti's locker [malware] became highly detectable and was rarely deployed. The only possible decision was to rebrand.

For over two months, Conti collective had been silently creating subdivisions that began operations before the start of the shutdown process. These subgroups either utilized existing Conti alter egos and locker malware, or took the opportunity to create new ones.

This decision was convenient for Conti, as they already had a couple of subsidiaries operating under different names: KaraKurt, BlackByte, BlackBasta. The rebranded version of Conti—the monster splitting into pieces still very much alive—ensured that whatever form Conti's ex-affiliates chose to take, they would emerge into the public eye **before** news of Conti's obsolescence could spread, controlling the narrative around the dissolution as well as significantly complicating any future threat attributions.

Their report then quotes Sun Tzu's "The Art of War", noting that:
"All warfare is based on deception. Hence, when we are able to attack, we must seem unable; when using our forces, we must appear inactive; when we are near, we must make the enemy believe we are far away; when far away, we must make him believe we are near."

And then they note that *"This is where the plans for what was left of Conti became increasingly complex."*

In order to hide the fact that Conti was now dispersed and operating via smaller, more novel brands, the former affiliates of the gang had to now convincingly simulate the actions of a dead brand.

Conti's remaining infrastructure operated like an army preparing for an ambush. Lingering actors were left to keep their fires lit, visible from behind enemy lines. Meanwhile, hidden from view, Conti's most skilled agents were instead laid low in a nearby encampment, biding their time while watching their great and empty camp send out smoke signals, meticulously emulating the movements of an active group.

Conti continued to publish documents stolen from victims (most likely targets hit earlier with attacks and lined up in a sort of "queue" for public release) and "campaigned" hard for themselves on criminal forums. Their public persona boasted a strong and enduring foundation, even one that was willing to further expand the group's operations. From the perspective of Conti's posting history, the group remained stronger than ever.

They then share a snapshot of a long and quite rambling chest-thumping post, from March 30th, where a Conti representative talks up the group's success, even seeking to recruit new affiliates.

However, in order to pull off their ultimate tactical maneuver, the agents left behind to operate from within Conti's massive, empty shell had to ensure that their antics would successfully lure attention away from their escaping comrades. To do this, they had to be certain that they left bait big enough to satisfy all of their opposing forces. Conti would have to perform a grand finale—one big enough to live up to the group's name.

And finally, on May 8th, Costa Rican President Rodrigo Chaves declared a national emergency as the result of a major cyber attack executed by the Conti ransomware gang. The massive attack, which took place against multiple Costa Rican government agencies, seems almost like a last-ditch effort by the group to squeeze a few more drops of riches from foreign government funds.

However, AdvIntel's unique adversarial visibility and intelligence findings led to, what was in fact, the opposite conclusion: The only goal Conti had for this final attack was to use the platform as a tool of publicity to perform their own death and subsequent rebirth.

AdvIntel has been tracking the preparations for this attack since April 14, 2022 — days before even the initial compromise. Our prevention alert was sent on April 15, 2022, three days before the first incident compromising Costa Rica's Ministry of Finance occurred.

Their report then links to a Tweet thread in spanish, but it appears to be dated from the 18th: https://twitter.com/HaciendaCR/status/1516190939114803203 But they then provide a screen shot which indeed appears to substantiate a three-day early warning of an impending attack.

So they explain:

In our pre-and-post attack investigation, we have found:

- The agenda to conduct the attack on Costa Rica for the purpose of publicity instead of ransom was declared internally by the Conti leadership.

- Internal communications between group members suggested that the requested ransom payment was far below $1 million USD (despite unverified claims of the ransom being $10 million USD, followed by Conti's own claims that the sum was $20 million USD). A low demand such as this, made to a state entity no less, was only made with the knowledge that the group would never see payment for the ransom.

- Conti was very vocal about the attack, constantly adding new political statements.

The attack on Costa Rica indeed brought Conti into the spotlight and helped them to maintain the illusion of life for just a bit longer, while the real restructuring had already taken place.

While Conti had been busy with its diversion tactics, other brands such as KaraKurt, BlackByte, and numerous other groups which existed as extensions of Conti, but without taking the group's name, were extremely operationally active, although they worked in silence.

Working concurrently with them, talented infiltration specialists, who were ultimately the backbone of Conti's gang, were also more active than ever, forming alliances with BlackCat, AvosLocker, HIVE, HelloKitty/FiveHands, and a whole other cadre of ransomware groups.

These pen-testers maintain personal loyalty to the people who created Conti but ultimately continued their work with other gangs in order to fully shed Conti's name and image.
This situation presents the first, and foremost reason for Conti's timely end—toxic branding.

Indeed, the first two months of 2022 left a major mark on the Conti name. While there is no tangible evidence to suggest that the well-known Conti leaks had any impact on the group's operations, the event which provoked the leak— Conti's claim to support the Russian government, seems to have been the fatal blow for the group, despite being revoked almost immediately.

We noted this at the time. What Conti posted was:

The statement had several key consequences, all of which deeply reshaped the environment Conti was operating within.

First, by engaging in political discourse, Conti broke the first unspoken rule of the Russian-speaking cybercrime community—not to intervene in state matters.

In AdvIntel's public blog regarding REvil's ultimate takedown by the Russian government, AdvIntel provided an in-depth analysis of this unspoken agreement, making case studies of the two most notable groups to break it—Avaddon and REvil. With the ongoing Russian invasion of Ukraine, it may be very plausible that Russia's state security apparatus is attempting to exert governmental control over its cyberspace, even taking down groups that appear to have been allies, but who exhibited undue independence with their actions.

AdvIntel had seen internal communication of the Conti leadership suggesting that the Russian FSB had been pressuring the group, and even though non-factual evidence was involved, the REvil scenario may have simply repeated itself with Conti, the group's brand becoming a target for Russian authorities despite their pledged loyalties.

Second, Conti's allegiance to the Russian invasion of Ukraine provoked internal conflict, and brought shame on the Conti name from members who were either ethnically Ukrainian, or were Russian but supported Ukraine, or simply wanted to maintain an anti-war ethic.

Considering that one of these members decided to betray the gang and leak private Conti chat logs not long after this conflict began, illustrated the final nail in Conti's self-made coffin:

The third, and most important factor—by pledging their allegiance to the Russian government, Conti as a brand became associated with the Russian state—a state that is currently undergoing extreme sanctions.

In the eyes of the state, each ransom payment going to Conti may have potentially gone to an individual under sanction, turning simple data extortion into a violation of OFAC regulation and sanction policies against Russia. This liability came to a head on May 6, 2022, when the US State Department offered rewards up to $10 million USD for information that led to the takedown of the Conti group.

As a result of these limitations, Conti had essentially cut itself off from the main source of income. Our sensitive source intelligence shows that many victims were prohibited from paying ransom to Conti. Other victims and companies who would have negotiated ransomware payments were more ready to risk the financial damage of not paying the ransom than they were to make payments to a pro-Russian state-sanctioned entity.

As AdvIntel previously stated, the end of the Conti brand does not equal the end of Conti as an organization. As seen with the Costa Rica case, Conti has been carefully planning its rebranding for several months, preparing a comprehensive strategy to execute it. This strategy is based on two pillars:

First, Conti is adopting a network organizational structure, more horizontal and decentralized than the previously rigid Conti hierarchy. This structure will be a coalition of several equal subdivisions, some of which will be independent, and some existing within another ransomware collective. However, they will all be united by internal loyalty to both each other and the Conti leadership, especially “reshaev”.

At this point, this network includes the following groups:

The first type being fully autonomous: No locker involved, pure data-stealing:

- Karakurt
- BlackBasta
- BlackByte

The second type being semi-autonomous: Acting as Conti-loyal collective affiliates within other collectives in order to use their locker:

- AlphV/BlackCat
- HIVE
- HelloKitty/FiveHands
- AvosLocker

The third type being Independent affiliates: Working individually, but keeping their loyalty to the organization.

And finally the fourth type being mergers & acquisitions where Conti leadership infiltrates a pre-existing minor brand and consumes it entirely, keeping the small brand name. The small group's leader loses their independence, but receives a massive influx of manpower, while Conti obtains a new subsidiary group.

This is different from Ransomware-as-a-Service, since this network, at least at the time of writing, does not seem to be accepting new members as part of its structure. Moreover, unlike RaaS, this model seems to value operations being executed in an organized, team-led manner. Finally, unlike RaaS, all the members know each other very well personally and are able to leverage these personal connections and the loyalty they bring.

This model is more flexible and adaptive than the previous Conti hierarchy while also being more secure and resilient than RaaS.

The other major development for this new ransomware model is the transition from data encryption to data exfiltration, covered extensively by AdvIntel in our analysis of Karakurt and BlackByte. In a nutshell, relying on pure data exfiltration maintains most major benefits of a data encryption operation, while avoiding the issues of a locker altogether. Most likely, this will become the most important outcome of Conti's re-brand.

The actors that formed and worked under the Conti name have not, and will no, cease their forward movement within the threat landscape—their impact will simply leave a different shape

So, to our listeners... if anyone in your cyber sphere announces that Conti has shut down and disbanded. Now we all know better.